Evolution of a PDS Jukebox Architecture

Gary N. Walker

Presented at:
Science Information Systems Interoperability Conference
8 November 1995

JPL
Jet Propulsion Laboratory
California Institute of Technology

Evolution of a PDS Jukebox Architecture

Gary N. Walker
Jet Propulsion Laboratory

This paper describes the decisions, issues, and evolution of CD-ROM jukebox deployment in the Planetary Data System.

The PDS Distribution Challenge

The PDSdistribution challenge is to distribute data to a large customer-base, quickly, at an acceptably, low cost.

Long ago, we decided to move away from tapes to reduce cost and increase longevity. In the process, we selected CD-ROM. We have avoided the enormous costs of maintaining a large operations staff and shipping department.

Our CD collection numbers about 475 titles and is growing at a rate of about 230 titles per year. To date, approximately 150,000 replicas have been produced. We contract with a vendor to master our discs and make approximately 300 replicas of each disc. The number of replicas varies by the data set. The sum of all mastering and replication costs is approaching $600K.

Distribution of CD replicas was not a mistake. CD-ROMs in users' hands have been the best way to do business. PDS users have indicated that PDS peer-reviewed CD-ROMs are valuable enablers in their scientific research. Now that users are connected to networks and disk drive costs have dropped, the average user is ready for electronic access and delivery. Budget pressures are compelling us to reevaluate which volumes we should distribute on plastic vs. electronically.

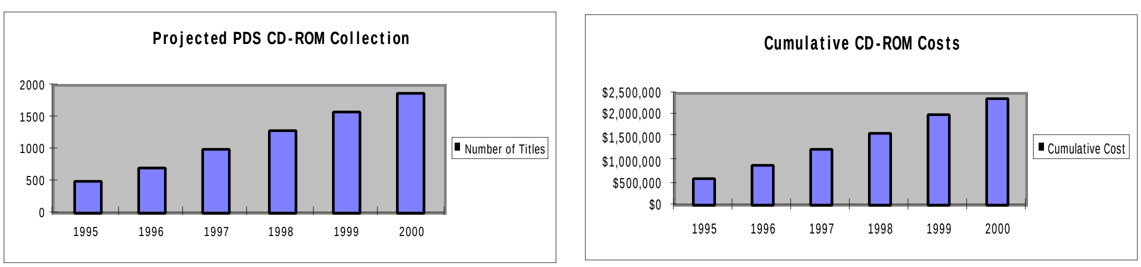

Figure 1
Collection and Cost Trends

Note from Figure 1 that we have a linear cost curve that is directly proportional to the number of titles produced. We have an advantageous contract with the CD vendor. There is no significant savings that could be gained by lowering vendor costs. There is no realistic way to change the relationship of the two curves if we continue to replicate and distribute CD-ROMs as in the past.

Why CD Jukeboxes?

The CD-ROM is the standard or base media type for PDS archived data. We selected CD-ROM, technically, for its portability and longevity and non-technically for its cost and size.

ISO 9660 is the CD-ROM standard. All PDS discs are written to the standard and are readable on just about any computer system and operating system that is in use today. Tapes were limited by their logical formats, e.g. ANSI, TAR, etc. and their physical formats e.g. 9-track, 7-track, 8mm, 4mm, QIC, etc.

CD-ROM life expectancy is greater than tape. CD-R has been tested and rated at 25-100 year and CD-ROM to 100-300 year expectancies.

Cost Issues

The graph on the right of Figure 2 compares the cost of placing a CD-R in a jukebox for each title versus mastering, replicating, and distributing CD-ROMs for each title. Note that the cost curves cross at about 26 titles.

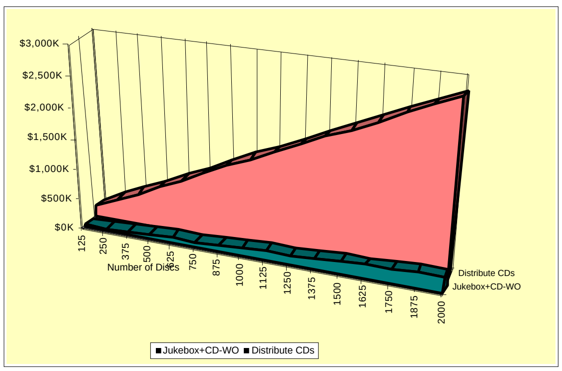

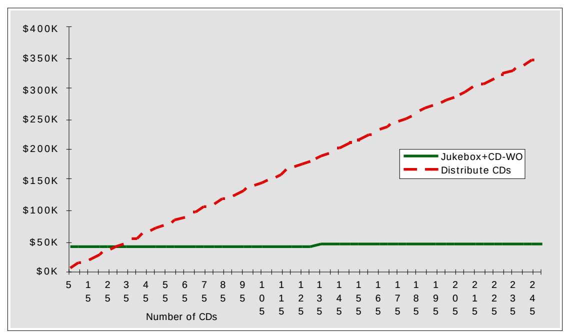

Figure 2
CD distribution vs. CD Jukeboxes

The graph on the left of Figure 2 compares the costs of mastering, replicating, and distributing up to 2000 titles versus the costs of setting up servers/workstations with (2) 250 disc jukeboxes. The server and jukebox method includes the costs of two CD-Rs and 0.125 FTE for system administration.

The two distribution methods displayed in the left chart are opposite ends of the cost spectrum. We feel a need to place PDS strategy somewhere in-between the poles. Some data needs to be distributed on disc. Large image files are classic examples. Giving users discs and the assurance associated with having a handy data set on the shelf are important. File type is a determining factor. Some data does not have wide interest; not many hobbyists need magnetometer or electric-field data.

Cost Issues

Although the per megabyte Winchester-technology disk drive cost is low, the cost of using disk drives is high when large data volumes must be supported. Disk drives have reliability and maintenance issues such as disastrous head-crashes; CD-ROM does not.

PDS does not maintain all of its data and metadata on disk farms because of cost and reliability issues. PDS disk farms would require 310Gb in 1995, 460GB in 1996, and would grow to 1.2 TB in 2000. Disk farms that large would require controllers and redundant arrays of inexpensive disk (RAID). RAID is not inexpensive. Disk sells for about 26¢ per MB which means that our current data-on-disk requirement would cost $81K today and would grow to $312K in 2000 AD. These expense figures do not include the costs for maintenance, controllers/RAID controllers, operations, facility space, and utilities.

The tradeoff is performance versus cost. CD-ROM drives have slower access times than disk 150 ms versus 10 ms. CD-ROM throughput is lower than disk throughput 400 KB/sec versus 10,000 KB/sec.

Device	Average Access	Throughput
1X CD-ROM	310ms	130KB/sec
2X CD-ROM	200ms	290KB/sec
4X CD-ROM	150ms	400KB/sec
6X CD-ROM	100ms	600KB/sec
Winchester-Tech. Drive	10ms	10,000KB/sec
Magneto-Optical	17ms	6,000KB/sec

Table 1
Typical Storage Device Specifications

Magneto-optical falls in between the hard disk and CD-ROM performance values, but the standards and physical media are changing too quickly to adopt as a PDS standard.

Technical Issues

PDS files range from 3 KB to 3^{+} MB each. The 3 MB file transfers or requests for 650 MB discs will allocate the disc drives for long periods of time.

Multiple file downloads will tie up a jukebox. You can *"bet the farm"* that some users will try to download entire CD volumes. Theoretically, a 3 MB file requires more than 9 minutes to download over 56 kbs links. This assumes the transfer can use the full bandwidth of the link and that data packets can be streamed. The best transfer rate that I have observed in my office is 8 kilobytes per second. The link from my building is a T1 (1.544 megabits per second). Imagine the allocation time of a 650 MB transfer request.

Contention for the disc readers or CD-ROM platters affects users. Many discs are available to be read by the shared drives in a jukebox. A file transfer allocates the drive and CD-ROM platter, forcing other users to wait in a queue or to be refused service. Clearly the issue is two-fold, drives and platters are points of contention. If a user wants data from an unallocated disc, but no drives are free, there is a delay. If a user wants data from a disc that is in use, the existence an unallocated drive won't matter.

Network bandwidth is out of our control. We can't control the bandwidth at the source or the target networks, let alone throughout the transfer path. The user has no control over the path a data transfer will follow.

Bus bandwidth and CPU performance on host platforms are limiting factors. Jukeboxes are not standalone; they must be attached to workstations or servers. We need to be cognizant of the possible bottlenecks on heavily used servers.

File indexing is an issue. There are thousands to millions of files to track on each jukebox. Clearly, 650 MB CD-ROMs can hold many files. Compound this by the number of discs that each jukebox can hold. It will take efficient indexing to allow a system to know what data files are in each jukebox and where.

Many file formats exist, but the *"right"* format for one user may not be useful to other users. The PDS has a marvelous, flexible image format that was designed to accommodate mission imaging team concepts. Some users may need to change the formats to other formats for their preferred image processing tools.

Administrative Issues

The legality of distributing data over the net is an issue. There are export laws and designated countries rules that must be followed. Some supporting data has limited rights. Published papers often become the property of the journals. PDS has been given limited rights to place some journal articles on CD-ROMs for limited distribution. Electronic distribution could violate copyright laws.

Decisions must be made as to where discs reside and in how many sites. A data system must make these decisions wisely to ensure appropriate user support and to ensure that valuable jukebox landscape is not squandered.

The jukebox-based distribution scheme must allow for redundancy and fault tolerance. A system must be able to handle equipment and network outages gracefully.

Solution Technologies

There are many simple and inexpensive solutions to technology issues. The *Sixty-four Thousand Dollar* question, do we implement these solutions in expectation of the problems or do we wait for the problems to rear their ugly heads, needs to be addressed.

Off-hour delivery of large orders completed in non-primetime helps to ensure acceptable access to all. Delayed delivery may try the patience of the requesters. We have all become jaded. We want it now! We forget that we used to have to wait days, weeks, months, and even years, relying upon *SneakerNet* and the various parcel delivery services. Large requests are still better met by distribution of CD media.

A hard disk cache for most used or most recently used files or volumes can remedy slow disc performance, contention for discs, and contention for CD readers.

A planetary inventory at each site that serves as an index and roadmap to all PDS sites can track multiple data-provider sites. Intelligent agents are needed to keep the planetary inventory up-to-date, automatically. The planetary inventory and multiple data sources would improve performance and reduce contention. PDS could use multicast to specific servers listed in the planetary directory or broadcast to all PDS jukebox servers, creating competition to fill orders. The servers with the requested data could respond in a horse-race fashion to fill the order.

On-the-fly data compression could improve performance and reduce contention. A threshold could be set to compress all transfers that are greater than a specific number of MB.

To make the data more usable to requesters, user-selectable formats with on-the-fly conversion could be offered. This concept has inherent efficiency issues.

Software and Access Paradigms

FTP and Gopher servers are common access paradigms available to essentially all Internet users. There are few drawbacks to using FTP or Gopher servers for delivery.

Electronic mail can be used for delivery of files. Electronic mail can be better used to notify users regarding delayed deliveries. Electronic delivery employing SMTP and MIME enclosures as file transport mechanisms have proven to be unreliable and fraught with problems. Some mail gateways do strange things to good data. Not all user mail systems interpret the enclosures well.

Thumbnails, browse images, and full resolution browse images can improve user access to image and table archives. These formats can reduce the size of transfers by allowing the user to choose the right files remotely, rather than to download many files and cull out non-useful files locally.

Intelligent user interfaces provide intuitive interfaces that make the archive media and organization transparent to the user. The PDS Image Node has done a nice job of providing an intuitive graphical selection interface of a Venus image map-style finder and some Martian landscape. The URLs to view these examples are:

> *http://cdwings.jpl.nasa.gov/PDS/public/magellan/midrcd_query.html*
> *http://cdwings.jpl.nasa.gov/PDS/public/vikingl/vl_images.html*

Software and access paradigms will affect the users' perception of the value of service. Remember that users of computers often forget the past in their expectation of immediate response. Each user tends to feel that he, alone, is the only user.

Jukebox Specifications

CD Jukeboxes	Num. Discs	Num. Drives	Mechanism	Cost	Cost/ MB	Discs/ Drive	Num. Gigabytes	Load Time
Kodak	100	1	Robotic	$4,275	$.066	100	65,000	<7 sec.
Pioneer DRM5004X	500	4	Robotic	$22,000	$.067	125	325,000	<17 sec.
Kubik	240	4	Carousel	$16,000	$.102	60	156,000	≈12 sec.
NSM, MDI	150	4	Robotic	$12,000	$.123	37.5	97,500	<5 sec.
Pioneer DRM624X	6	1	1 Cart.	$625	$.16	6	3,900	<5 sec.
Pioneer DRM1804X	18	1	3 Cart.	$2025	$.173	18	11,700	<18 sec.

Table 2
Typical Jukebox Specifications

Table 2 provides sample jukebox specifications for several common products. Please note that average and approximate numbers are presented. Some of the numbers are provided by the manufacturers, other information was obtained through research and the World-Wide Web. Check the manufacturers' specifications, other user experiences, and (better yet) test for yourself before you buy. Do not assume 100% accuracy of values in this chart.

The key point of this section is that there are performance, capacity, and cost issues that must be resolved by engineering personnel for each jukebox application. Disc quantity, drive quantity, disc load/unload, and physical disc addition/deletion are important. The speed of the individual drive is less important, as manufacturers are constantly improving their hardware.

Driver software is an issue. Is it supplied with the drive, is it purchased elsewhere, or do you grow your own? Our experience with the latter option has not been good and we recommend purchasing commercial drivers.

The PDS Jukebox Evolution

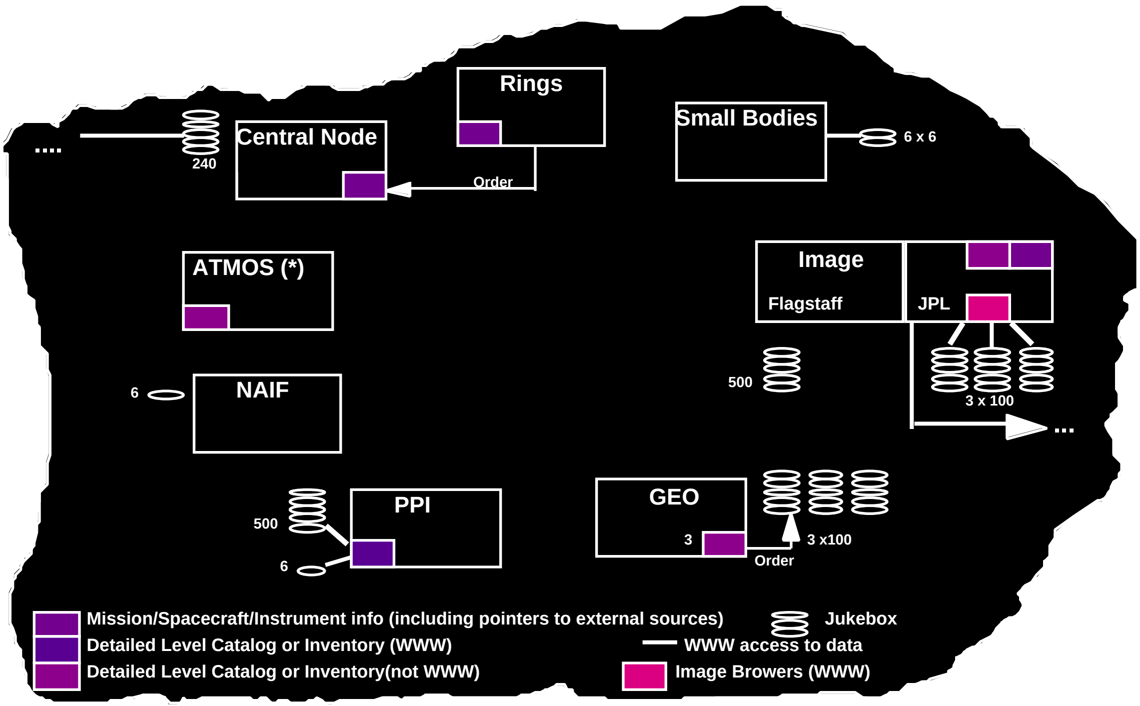

Figure 3
PDS Systems in 1995

Figure 3 shows how PDS has deployed its jukeboxes and other on-line services. The jukebox deployment has been based upon node budgets and node needs. Each node operates in a method that is determined by its discipline community and PDS system requirements.

There has been system engineering at the node levels to employ jukeboxes to meet discipline science community needs. The PDS has begun reengineering to coordinate future jukebox growth and current resource redeployment across the whole system.

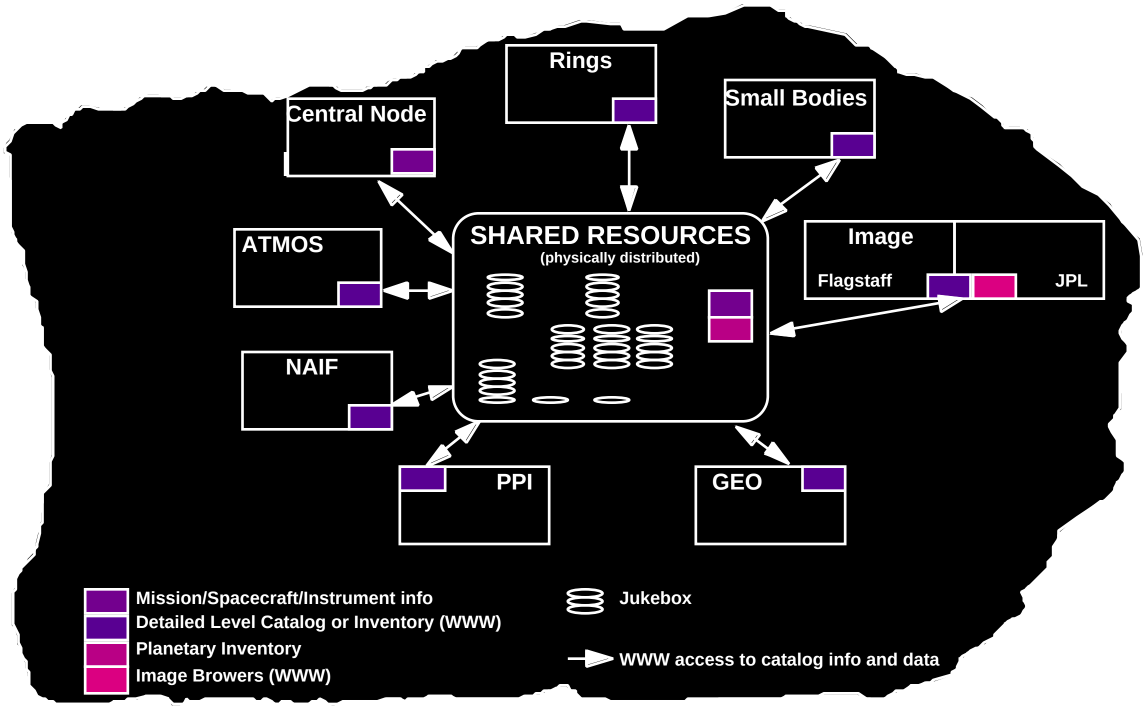

Figure 4
PDS Systems, Future Target

Our future target is pictured In Figure 4. All jukeboxes within the PDS structure will be coordinated and distributed with replicated disc deployment. The locations will be transparent to the user. The user will see all that is applicable to satisfy a request. Multiple user interface strategies will be allowed and a common access mechanism will be available.

The shared resources in Figure 4 will be distributed. The figure does not denote location of the jukebox resources.

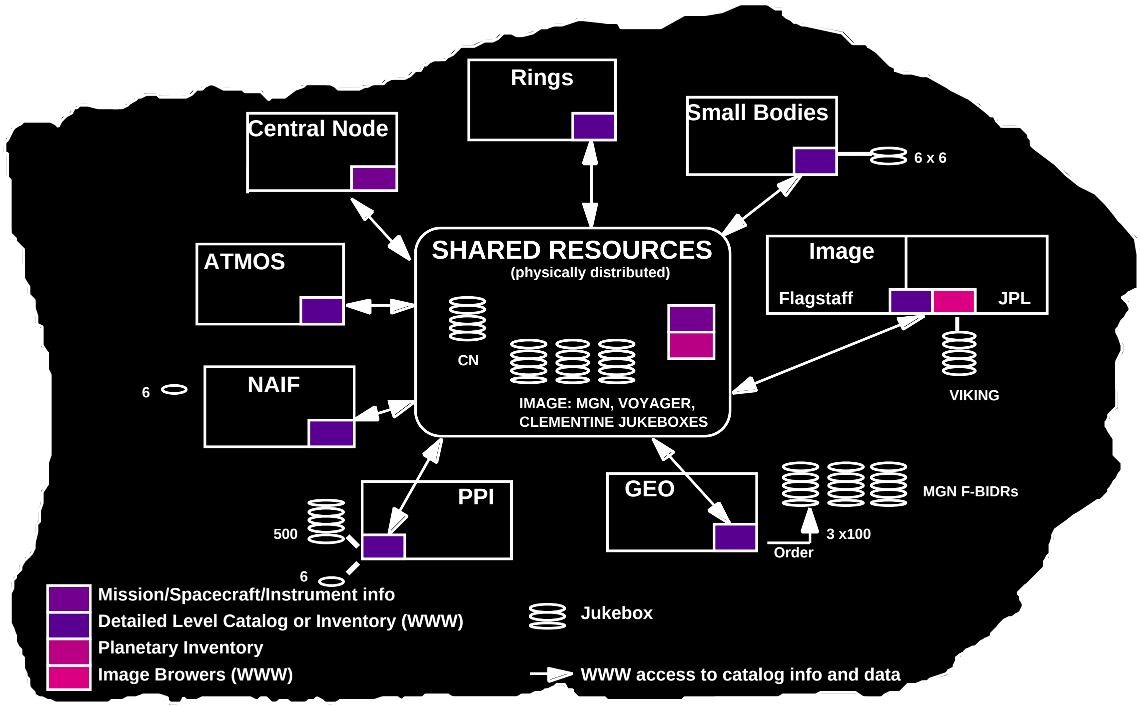

Figure 5
PDS Systems, 1996 Plant

Figure 5 shows our jukebox deployment plans for 1996. During 1996 the PDS will begin to coordinate and redeploy its jukebox resources. We have a pilot program called *The PDS Navigator* in process now.

The *Navigator* pilot includes three nodes. The Image Co-nodes, Geosciences Node, and Central Node are participating in the prototype activity using Clementine data. Users may search for Clementine files via the user interfaces at any of the sites. The user will be notified by e-mail when the data is staged at an ftp site. The user is allotted a specific time duration during which the data may be downloaded. The data is distributed among jukeboxes and replicated at different sites. The user will not see the locations when searching for data.

Conclusion

The Planetary Data System has been a pathfinder in data distribution. The PDS is constantly reevaluating technology, cost, and user needs to try to provide the greatest amount of data to the widest possible audience at the most appropriate cost.

We moved from tape to CD-ROM based upon needs, technology, and cost. We are now modifying our highly successful physical distribution program to include electronic distribution of our massive planetary archives. The electronic distribution will make NASA's planetary treasures accessible to the science community as well as the American public.

We are encountering administrative and technology challenges as we evolve. We are meeting these challenges with sound engineering.

Bibliography

Anderson, Terri. PDS System Document. August 1995.

Arvidson, R. The PDS Navigator. PDS Whitepaper. July 1995.

Comer, Douglas. Internetworking With TCP/IP Principles, Protocols, and Architecture. Prentice Hall, 1988

Data Distribution Laboratory. Producing Multiplatform CD-ROMs. Science Information Systems Newsletter, Issue 27, January 1993.

Martin, Mike. Compact Disc Media Evaluation, Study Shows That The Plastic Disks Will Not Self-Destruct. Science Information Systems Newsletter, Issue 26, October 1992.

Martin, M., Rengarajan, K, Hughes, S., and Martinez, A. Adventures With Jukeboxes. Science Information Systems Newsletter, Issue 28, April 1993.

Martin, M. And Rengarajan, K. Adventures With Jukeboxes, Part 2. Science Information Systems Newsletter, Issue 29, July 1993.

Plain Talk: CD-ROM Lifetime. 3M Optical Recording. 1992.